SUPPLEMENTARY INFORMATION

Accounting for saccade latencies

Our task was not optimized to detect reaction times, as there was a forced delay after both options and then a required fixation before the saccade was permitted. Saccade delays were thus not reaction times in the traditional sense, and likely did not reflect the decision process. A concern one might have is that these response times, nonetheless, had some effect on firing rates that interfered with the effects we report here. To determine whether this was the case, we replicated all of our single-unit modulation analyses while including response times in the regression. In all but two cases, our results remain unchanged qualitatively. We highlight exceptions below in bold. (1) In sgACC, encoding for chosen value becomes significant (suggesting that waiting time is not a confound but a source of noise). (2) Encoding of number of tokens in the reward epoch is no longer significant. We reproduce these numbers below, and add them to the supplementary information section (all statistical tests were conducted the same as in the original manuscript; fractions were tested using a two-sided binomial test at alpha = 0.05, and modulation bias in the overall population was quantified using a Wilcoxon signed rank test).

In dACC:

Variable	% significantly modulated	% positively modulated	Modulation bias *Z-stat (p-value)*
Offer 1 value (attended)	24.0% (31/129, p < 0.0001)	51.6% (16/31, p = 1.00)	-0.340 (p = 0.734)
Offer 1 value (remembered)	20.9% (27/129, p < 0.0001)	55.6% (15/27, p = 0.701)	-0.963 (p = 0.336)
Offer 2 value (attended)	14.0% (18/129, p = 0.0001)	61.1% (11/18, p = 0.481)	0.384 (p = 0.701)
Outcome	27.1% (35/129, p < 0.0001)	68.6% (24/35, p = 0.0410)	1.82 (p = 0.0694)
Number of tokens	30.2% (39/129, p < 0.0001)	51.3% (20/39, p = 1.00)	0.384 (p = 0.701)
Jackpot	38.8% (50/129, p < 0.0001)	60.0% (30/50, p = 0.203)	1.30 (p = 0.195)

In sgACC:

Variable	% significantly modulated	% positively modulated	Modulation bias

			Z-stat (p-value)
Offer 1 value (attended)	11.0% (16/146, p = 0.0028)	68.8% (11/16, p = 0.210)	1.480 (p = 0.139)
Offer 1 value (remembered)	7.53% (11/146, p = 0.178)	45.5% (5/11, p = 1.00)	-0.646 (p = 0.519)
Offer 2 value (attended)	8.22% (12/146, p = 0.0843)	58.3% (7/12, p = 0.774)	0.962 (p = 0.336)
Outcome	14.4% (21/146, p < 0.0001)	19.1% (4/21, p = 0.0072)	-3.52 (p = 0.0004)
Number of tokens	**8.22% (12/146, p = 0.0843)**	25.0% (3/12, p = 0.146)	-2.47 (p = 0.0136)
Jackpot	11.0% (16/146, p = 0.0028)	87.5% (14/16, p = 0.0042)	4.098 (p < 0.0001)

Subgenual ACC subregions

Of the 69 sgACC cells recorded in subject J, 29 of these were recorded in area 32, and 40 were recorded in area 25. All 77 neurons recorded in subject B were recorded in area 25. Given that only ~20% of our neurons were recorded in area 32, we do not think this fraction is sufficient to detect functional differences.

We wanted to know to what degree neurons in area 32 account for the task-relevant responses we have reported. To that end, we quantified the fraction of significantly-modulated cells for each variable that were recorded in area 32 vs. area 25, and compared this to the total ratio of neurons recorded from each area. Our logic was that, if these fractions are significantly different, our results would be more representative of one subregion. Specifically, it is possible that very strong tuning in area 32 accounts for all the significantly-modulated neurons we have observed, and area 25 neurons merely add noise. As shown in the table below, this was not the case: neurons in area 32 accounted for a similar fraction of modulation as would be expected given their frequency in the population. We thus have no strong reason to conclude that functionality between these areas was different in our task.

Variable	Fraction of significantly-modulated population recorded in area 32	Comparison of proportions (in significant population vs. in overall population, using a Z-score test of proportions) *Z statistic (p-value)*
Offer 1 value (attended)	2/16	0.711 (p = 0.478)
Offer 1 value	1/11	0.876 (p = 0.379)

(remembered)		
Offer 2 value (attended)	1/12	0.979 (p = 0.327)
Outcome	1/22	1.75 (p = 0.0801)
Jackpot	3/18	0.323 (p = 0.749)
Number of tokens	2/14	0.504 (p = 0.617)

Example mutual inhibition cells

Figure S1 shows two neurons, one from dACC and one from sgACC, that exhibit mutual inhibition. These neurons both fire more vigorously when offer 1 is lower than offer 2 in value. Neurons may also show the opposite response. Note that mutual inhibition is a *population* signature of value comparison, and may exist at the level of the population but not be instantiated in any single neuron. However, in our data, neurons encoding the two offer values overlap significantly (see main text).

Figure S1: *Neural responses to difference in offer values. Grey, shaded region indicates the time when offer 2 was displayed on the screen and offer 1 was presumably held in working memory (i.e. epoch of comparison).* ***A:*** *Responses of one dACC cell whose firing rate increased in response values of offer 2 larger than offer 1.* ***B:*** *Responses of one sgACC cell whose firing rate increased in response to values of offer 2 larger than offer 1. See main text for population statistics.*